arXiv:1404.2670v1 [math.NA] 10 Apr 2014

SCALABLE AND FAULT TOLERANT COMPUTATION WITH THE SPARSE GRID COMBINATION TECHNIQUE*

BRENDAN HARDING†, MARKUS HEGLAND†, JAY LARSON†, AND JAMES SOUTHERN‡

Abstract. This paper continues to develop a fault tolerant extension of the sparse grid combination technique recently proposed in [B. Harding and M. Hegland, *ANZIAM J.*, 54 (CTAC2012), pp. C394–C411]. The approach is novel for two reasons, first it provides several levels in which one can exploit parallelism leading towards massively parallel implementations, and second, it provides algorithm-based fault tolerance so that solutions can still be recovered if failures occur during computation. We present a generalisation of the combination technique from which the fault tolerant algorithm is a consequence. Using a model for the time between faults on each node of a high performance computer we provide bounds on the expected error for interpolation with this algorithm. Numerical experiments on the scalar advection PDE demonstrate that the algorithm is resilient to faults on a real application. It is observed that the trade-off of recovery time to decreased accuracy of the solution is suitably small. A comparison with traditional checkpoint-restart methods applied to the combination technique show that our approach is highly scalable with respect to the number of faults.

Key words. exascale computing, algorithm-based fault tolerance, sparse grid combination technique, parallel algorithms

AMS subject classifications. 65Y05, 68W10

1. Introduction.

Many recent survey articles on the challenges of achieving exascale computing identify three issues to be overcome: exploiting massive parallelism, reducing energy usage and, in particular, coping with run-time failures [3, 6, 2]. Faults are an issue at peta/exa-scale due to the increasing number of components in such systems. Traditional checkpoint-restart based solutions become unfeasible at this scale as the decreasing mean time between failures approaches the time required to checkpoint and restart an application. Algorithm based fault tolerance has been studied as a promising solution to this issue for many problems [14, 1].

Sparse grids were introduced in the study of high dimensional problems as a way to reduce the *curse of dimensionality*. They are based on the observation that when a solution on a regular grid is decomposed into its hierarchical bases the highest frequency components contribute the least to sufficiently smooth solutions. Removing some of these high frequency components has a small impact on the accuracy of the solution whilst significantly reducing the computational complexity [8, 7]. The combination technique was introduced to approximate sparse grid solutions without the complications of computing with a hierarchical basis. In recent years these approaches have been applied to a wide variety of applications from real time visualisation of complex datasets to solving high dimensional problems that were previously cumbersome [19, 8].

Previously [9, 18, 10] it has been described how the combination technique can be implemented within a *Map Reduce* framework. Doing so allows one to exploit an extra layer of parallelism and fault tolerance can be achieved by recomputing

*This research was supported by the Australian Research Council's *Linkage Projects* funding scheme (project number LP110200410). We are grateful to Fujitsu Laboratories of Europe for providing funding as the collaborative partner in this project.

†Mathematical Sciences Institute, Australian National University, Canberra, Australian Capital Territory, Australia, 0200.

‡Fujitsu Laboratories of Europe, Hayes Park Central, Hayes End Road, Hayes, Middlesex, UB4 8FE, United Kingdom.

failed map tasks as described in [4]. Also proposed was an alternative approach to fault tolerance in which recomputation can be avoided for a small trade off in solution error. In [11] we demonstrated this approach for a simple two-dimensional problem showing that the average solution error after simulated faults was generally close to that without faults. In this paper we develop and discuss this approach in much greater detail. In particular we develop a general theory for computing new combination coefficients and discuss a three-dimensional implementation based on MPI and OpenMP which scales well for relatively small problems. As has been done in the previous literature [9, 18, 10], we use the solution of the scalar advection PDE for our numerical experiments.

The remainder of the paper is organised as follows. In Section 2 we review the combination technique and provide some well-known results which are relevant to our analysis of the fault tolerant combination technique. We then develop the notion of a general combination technique.

In Section 3 we describe how the combination technique can be modified to be fault tolerant as an application of the general combination technique. Using a simple model for faults on each node of a supercomputer we are able to model the failure of component grids in the combination technique and apply this to the simulation of faults in our code. We present bounds on the expected error and discuss in how faults affect the scalability of the algorithm as a whole.

In Section 4 we describe the details of our implementation. In particular we discuss the multi-layered approach and the way in which multiple components work together in order to harness the many levels of parallelism. We also discuss the scalability bottleneck caused by communications and several ways in which one may address this.

Finally, in Section 5 we present numerical results obtained by running our implementation with simulated faults on a PDE solver. We demonstrate that our approach scales well to a large number of faults and has a relatively small impact on the solution error.

2. The Combination Technique and a Generalisation.

We introduce the combination technique and a classical result which will be used in our analysis of the fault tolerant algorithm. For a complete introduction of the combination technique one should refer to [5, 8, 7]. We then go on to extend this to a more general notion of a combination technique building on existing work on adaptive sparse grids [12].

2.1. The Combination Technique.

Let $i \in \mathbb{N}$, then we define $\Omega_i := \{k2^{-i} : k = 0, \dots, 2^i\}$ to be a discretisation of the unit interval. Similarly for $i \in \mathbb{N}^d$ we define $\Omega_i := \Omega_{i_1} \times \cdots \times \Omega_{i_d}$ as a grid on the unit d-cube. Throughout the rest of this paper we treat the variables i, j as multi-indices in $\mathbb{N}^d$. We say $i \leq j$ if and only if $i_k \leq j_k$ for all $k \in \{1, \dots, d\}$, and similarly, $i < j$ if and only if $i \leq j$ and $i \neq j$.

Now suppose we have a problem with solution $u \in V \subset C([0,1]^d)$, then we use $V_i \subset V$ to denote the function space consisting of piecewise linear functions uniquely determined by their values on the grid Ω_i. Further, we denote an approximation of u in the space V_i by u_i. The sparse grid space of level n is defined to be $V_n^s := \sum_{\|i\|_1 \leq n} V_i$. A sparse grid solution is a $u_n^s \in V_n^s$ which closely approximates $u \in V$. The combination technique approximates a sparse grid solution by taking the sum of several solutions from different anisotropic grids. The classical combination technique

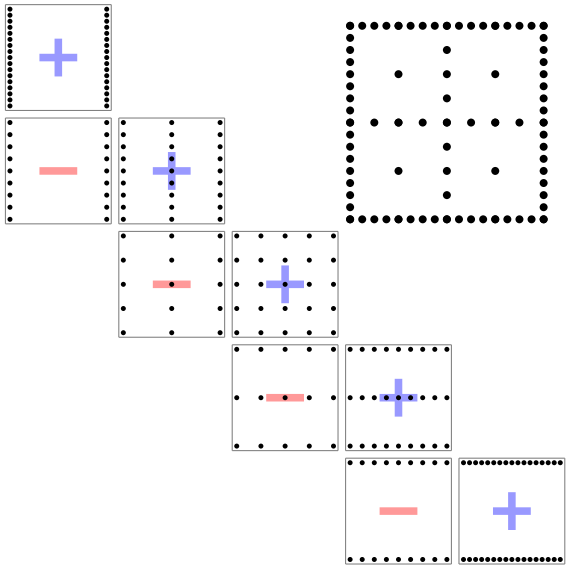

FIG. 2.1. *A level 4 combination in 2 dimensions. The 9 component grids are arranged according to the frequency of data points in each dimension. A plus or minus denotes a combination coefficient of +1 or −1 respectively. In the top right is the (enlarged) sparse grid corresponding to the union of the component grids.*

is given by the equation

$$u_n^c := \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} \sum_{\|i\|_1 = n-k} u_i \,. \tag{2.1}$$

Fundamentally, this is an application of the inclusion/exclusion principle. This can be seen if the function spaces are viewed as a lattice [12]. For example, if one wishes to add the functions $u_i \in V_i$ and $u_j \in V_j$ then the result will have two contributions from the intersection space $V_{i \wedge j} = V_i \cap V_j$, with $i \wedge j = (\min\{i_1, j_1\}, \ldots, \min\{i_d, j_d\})$. To avoid this we simply take $u_i + u_j - u_{i \wedge j}$. This can be seen in Figure 2.1, which shows a level 4 combination in 2 dimensions.

An important concept in the development of sparse grids is that of the hierarchical space (sometimes also referred to as a hierarchical surplus). A simple definition of the hierarchical space W_i is the space of all functions $f_i \in V_i$ such that f_i is zero when sampled on all grid points in the set $\bigcup_{j<i} \Omega_j$. Equivalently we have $V_i = W_i \oplus \sum_{j<i} V_j$. Noting that $V_i = \bigoplus_{j \leq i} W_j$, a hierarchical decomposition of $u_i \in V_i$ is the computation of the unique components $h_j \in W_j$ for $j \leq i$ such that $u_i = \sum_{j \leq i} h_j$. The sparse grid space can also be written in terms of hierarchical spaces as $V_n^s = \bigoplus_{\|i\|_1 \leq n} W_i$.

Let H^2_{mix} be the Sobolev-space with dominating mixed derivatives with norm

$$\|u\|^2_{H^2_{\text{mix}}} = \sum_{\|i\|_\infty \leq 2} \left\| \frac{\partial^{\|i\|_1}}{\partial x^i} u \right\|_2^2 \tag{2.2}$$

If $u \in H^2_{\text{mix}}$ then we have the estimate $\|h_j\|_2 \leq 3^{-d} 2^{-2\|i\|_1} |u|_{H^2_{\text{mix}}}$ for each of the hierarchical spaces where $|u|_{H^2_{\text{mix}}} := \left\| \frac{\partial^{2d}}{\partial x_1^2 \cdots \partial x_d^2} u \right\|_2^2$ is a semi-norm [5]. In the classical

theory of the combination technique this estimate is used to prove the error bound

$$\|u - u_n^c\|_2 \leq \sum_{\|j\|_1 > n} \|h_j\|_2 \leq 3^{-d}|u|_{H^2_{\text{mix}}} \sum_{k=n+1}^{\infty} 2^{-2k}\binom{k+d-1}{d-1} \tag{2.3}$$

$$= \frac{1}{3}\cdot 3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}} \sum_{k=0}^{d-1}\binom{n+d}{k}\left(\frac{1}{3}\right)^{d-1-k} \tag{2.4}$$

$$= \frac{1}{3}\cdot 3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}} \left(\frac{n^{d-1}}{(d-1)!} + \mathcal{O}(n^{d-2})\right) \tag{2.5}$$

for sparse grid interpolation. Similar bounds can be shown for the ∞ and energy norms [7].

In practice, strongly anisotropic grids, i.e. those with $\|i\|_\infty \approx \|i\|_1$, can be problematic. Not only are they difficult to compute in some circumstances but one also finds that they give poor approximations that do not cancel out in the combination as expected. Therefore it is often beneficial to implement a *truncated* combination. In this paper we define a truncated combination as

$$u^c_{n,\tau} := \sum_{k=0}^{d-1}(-1)^k\binom{d-1}{k} \sum_{\substack{\|i\|_1 = n-k \\ \min(i)\geq\tau}} u_i\,, \tag{2.6}$$

where τ is referred to as the truncation parameter. Such combinations will be used for the numerical results presented in Section 5.

2.2. A General Combination Technique. The combination technique can be generalised into arbitrary sums of solutions. Given a (finite) set of multi-indices $I \subset \mathbb{N}^d$, one can write

$$u^c_I = \sum_{i\in I} c_i u_i\,, \tag{2.7}$$

where the c_i are referred to as the combination coefficients. It is easy to see that $u^c_I \in V^s_I := \sum_{i\in I} V_i = \bigoplus_{i\in I\downarrow} W_i$ where $I\downarrow = \{i \in \mathbb{N}^d : \exists j \in I \text{ s.t. } i \leq j\}$. Not every choice of the c_i will produce a reasonable approximation to u. The question now is which combination coefficients produce the best approximation to u? One could attempt to solve the optimisation problem of minimising for example $\|u - \sum_{i\in I} c_i u_i\|_2$ which is known as *Opticom* [13]. However this is cumbersome to solve in a massively parallel implementation and also requires an approximation of the residual. In this paper we consider combinations which we know *a priori* will give a good approximation in some sense. We define a *sensible* combination to be one in which each hierarchical space contributes either once to the solution, or not at all. In particular we would like the coefficients to follow an inclusion/exclusion as described in Section 2.1. We observe that $W_i \subset V_j$ for all $j \geq i$, therefore one can easily determine how many times a given W_i contributes to the solution by summing all of the coefficients c_j for which $W_i \subset V_j$. In light of this we have the following definition.

DEFINITION 2.1. *A set of combination coefficients $\{c_i\}_{i\in I}$ is said to be* ***valid*** *if for each $i \in I$ it satisfies the property*

$$\sum_{j\in I,\, j\geq i} c_j \in \{0, 1\}\,.$$

We will also refer to a combination (of solutions) as valid if the corresponding set of combination coefficients are valid.

Another common assumption is that $\sum_{j\in I} c_j = 1$. This relates to the definition if $\underline{0} \in I$ in which case we could assert that $1 = \sum_{j\in I,\, j\geq \underline{0}} c_j$. The motivation for the definition is that this property is satisfied for dimension adaptive sparse grids [12]. Let $P_i : V \mapsto V_i$ be a lattice of projection operators associated with the tensor product space V. As in P_i satisfies $P_iP_j = P_{i\wedge j}$ (where $(i \wedge j)_k = \min\{i_k, j_k\}$), $P_iP_j = P_jP_i$ and $P_iP_i = P_i$. Defining $P_I : V \mapsto V_I^s$ it follows from [12, p. C344] that for $i \in I$

$$P_iP_I = P_i(1 - \prod_{j\in I}(1 - P_j)) = P_i - P_i(1 - P_i) \prod_{j\in I\setminus\{i\}} (1 - P_j) = P_i \,.$$

Since $P_I = \sum_{i\in I} c_iP_i$ it follows that $\sum_{\{j\in I \text{ s.t. } j\wedge i=i\}} c_i = \sum_{j\in I,\, j\geq i} c_i = 1$.

For a given I there are a many sets of valid combination coefficients that one might take. To determine which of these *a priori* will give the best approximation of u we use approximations which are based on sparse grid interpolation. It is reasonable to expect that such combinations will work well in a more general setting given the underlying inclusion/exclusion principle. The error estimate (2.3) for sparse grid interpolation is extended to the general combination technique by

$$\|u - u_I^c\|_2 \leq 3^{-d}|u|_{H^2_{\text{mix}}} \sum_{i\in\mathbb{N}^d} \left(4^{-\|i\|_1} \left| 1 - \sum_{j\in I,\, j\geq i} c_j \right| \right) .$$

Finding a set of valid coefficients which minimises this this bound is then equivalent to maximising

$$Q(\{c_i\}_{i\in I}) := \sum_{i\in I\downarrow} 4^{-\|i\|_1} \sum_{j\in I,\, j\geq i} c_j \,. \tag{2.8}$$

Therefore, the general problem of finding the best combination of solutions u_i for $i \in I$ can be formulated as an optimisation problem, in particular the maximisation of $Q(\{c_i\}_{i\in I})$ subject to the constraints $\sum_{j\in I,\, j\geq i} c_j \in \{0, 1\}$ for each $i \in I$. Since the c_i must be integers this would be a simple integer linear programming (ILP) problem if not for the non-trivial constraints. (To see why the c_i must be integers, we note that the non-zero c_i for which $c_j = 0$ for all $j > i$ must be 1 in order for the set of coefficients to be valid. The remaining coefficients are now obtained from the application of the inclusion/exclusion principle which can only result in integer coefficients.)

Fortunately we can simplify this by introducing the hierarchical coefficient.

DEFINITION 2.2. *Let I be a set of multi-indices, then for $i \in I\downarrow$ we define the* ***hierarchical coefficient***

$$w_i := \sum_{j\in I,\, j\geq i} c_j \,. \tag{2.9}$$

Suppose we expand our list of coefficients to the set $\{c_i\}_{i\in I\downarrow}$ with the assumption $c_i = 0$ for $i \notin I$. Now let c, w be vectors for the sets $\{c_i\}_{i\in I\downarrow}$, $\{w_i\}_{i\in I\downarrow}$, respectively (both having the same ordering with respect to $i \in I\downarrow$). Using (2.9) we can write $w = Mc$ where M is an $|I\downarrow| \times |I\downarrow|$ matrix. Further, we note that if the elements of c, w are ordered according to ascending or descending values of $\|i\|_1$ then M is an

upper or lower triangular matrix respectively with 1's on the diagonal. Therefore M is invertible and we have $c = M^{-1}w$. Additionally, the restriction that $c_i = 0$ for $i \notin I$ can be written as $(M^{-1}w)_i = 0$.

Since $w_i \in \{0, 1\}$ for any set of valid coefficients we can formulate the *general coefficient problem* (GCP) as the binary integer programming (BIP) problem of maximising

$$Q'(w) := \sum_{i \in I\downarrow} 4^{-\|i\|_1} w_i \tag{2.10}$$

with the equality constraints $(M^{-1}w)_i = 0$ for $i \notin I$. This is much more manageable in practice and can be solved using a variety of algorithms that are typically based on branch and bound, and/or cutting plane techniques. However, this formulation also reveals that the general coefficient problem is NP-complete [16]. Equivalently, one can also think of this as a weighted maximum satisfiability problem (Weighted MAX-SAT). If I is a downset then we can solve this rather quickly, but in general there exist cases which take an incredibly long time to solve. Another problem one runs into is that there is often not a unique solution. In such circumstances we will simply pick any one of the solutions as they cannot be further distinguished without additional information about u.

The only way to guarantee that a solution can be found quickly is to carefully choose the index set I. One particular class of index sets of interest are those which are closed under the $\wedge$ operator, that is if $i, j \in I$ then $i \wedge j \in I$. In the theory of partially ordered sets, $(I, \leq)$ with this property is referred to as a lower semi-lattice. For such I there is a unique solution to the GCP, namely $w_i = 1$ for all $i \in I\downarrow$ which clearly maximises (2.10). Computationally the coefficients can be found quickly by first finding $\max I := \{i \in I : \nexists j \in I \text{ s.t. } j > i\}$, setting $c_i = 1$ for $i \in \max I$, and then using the inclusion/exclusion principle to find the remaining coefficients in the order of descending $\|i\|_1$. This can also be viewed as an application of the lattice theory of projections on function spaces presented by Hegland [12].

Whilst the GCP presented is based upon $u \in H^2_{mix}$ we anticipate that the resulting combinations will still yield reasonable results for larger function spaces. This is based on the observation that the classical combination technique has been successfully applied to a wide variety of problems for which $u \notin H^2_{\text{mix}}$.

REMARK 1. *The restriction of w_i to binary variables can be relaxed to reals if we change the quantity we intend to optimise. The important observation to make here is that one would expect having two contributions from a hierarchical space is comparable to having no contributions. Further having a fractional contribution like $\frac{1}{2}$ would be better than having no contribution at all. In light of this we can try to solve the linear programming problem of minimising*

$$\sum_{i \in I\downarrow} 4^{-\|i\|_1} |1 - w_i| \tag{2.11}$$

subject to the equality constraints $(M^{-1}w)_i = 0$ for $i \notin I$. If I is closed under $\wedge$ then a minimum of 0 is achieved for the same hierarchical coefficients found in the binary formulation. In other cases the solution to this relaxed optimisation problem is no worse than the solution to the binary problem. This relaxation means the combination may no longer follow an inclusion/exclusion principle. In practice the results are highly dependant upon the true solution u and the approximation properties of each

of the u_i. Additionally, the non-differentiability of equation (2.11) *means that the problem is still non-trivial to solve in practice.*

An approach that can sometimes speed up the computation of a solution to this problem is to first find the solution to the quadratic programming problem of minimising

$$\sum_{i \in I\downarrow} 4^{-\|i\|_1}(1 - w_i)^2$$

subject to the same equality constraints. This is a linear problem which is easily solved using the method of Lagrange multipliers for example. In most circumstances we would expect the solution of this problem to be close to the minimum of equation (2.11) *and therefore make a good initial guess.*

3. Fault Tolerant Combination Technique and Fault Simulation.

3.1. Fault Tolerant Combination Technique (FTCT). In [9, 10] a fault tolerant combination technique was introduced. The most difficult aspect of generalising this work is the updating of coefficients. Whilst some theory and a few simple cases have been investigated, no general algorithm has been presented. Given the development of the general combination technique in Section 2.2 we are now able to consider a more complete theory of the FTCT.

Suppose we have a set I of multi-indices for which we intend to compute each of the solutions u_i and combine as in (2.7). As each of the u_i can be computed independently the computation of these is easily distributed across different nodes of a high performance computer. Suppose that one or more of these nodes experiences a fault, hardware or software in nature. As a result, some of our u_i may not have been computed correctly. We denote $J \subset I$ to be the set of indices for which the u_i where not correctly computed. A lossless approach to fault tolerance would be to recompute u_i for $i \in J$. However, since recomputation is often costly, we propose a lossy approach to fault tolerance in which the failed solutions are not recomputed. In this approach, rather than solving the generalised coefficient problem (GCP) for I, we instead solve it for $I \backslash J$. As $(I \backslash J)\downarrow \subseteq I\downarrow$ we expect this solution to have a larger error than that if no faults had occurred. However, if $|J|$ is relatively small we would also expect the loss of accuracy to be small because of the redundancy in the set of $\{u_i\}_{i \in I}$.

As discussed in Section 2.2, the GCP is difficult to solve in its most general form. Whilst it can be solved rather quickly if the poset $(I, \leq)$ is a lower semi-lattice, this is no longer any help in the FTCT since the random nature of faults means we cannot guarantee that $(I \backslash J, \leq)$ is always a lower semi-lattice. The only way we could ensure this is to restrict which elements of I can be in J. A simple way to achieve this is to recompute missing u_i if $(I \backslash \{i\}, \leq)$ is not a lower semi-lattice. In particular this is achieved if all u_i with $i \notin \max I$ are recomputed. Since elements in $\max I$ correspond to the solutions on the largest of the grids, we are avoiding the recomputation of the solutions which take the longest to compute. Additionally, this also means only the largest of the hierarchical spaces are ever omitted as a result of a failure. As these contribute the least to the solution we expect the resulting error to be relatively close to that of the solution if no faults had occurred. Finally, since $(I \backslash J, \leq)$ is then a lower semi-lattice, the resulting GCP for $I \backslash J$ has a unique maximal solution which is easily computed.

We now illustrate this approach as it is applied to the classical combination technique. We define $I_n = \{i \in \mathbb{N}^d : \|i\|_1 \leq n\}$. It was shown in [9] that the proportion

of additional unknowns in computing the solutions u_i for all $i \in I_n$ compared to $n-d < \|i\|_1 \leq n$ is at most $\frac{1}{2^d-1}$. If no faults occur then the combination is exactly the classical combination technique with $c_i = (-1)^{n-\|i\|_1}\binom{d-1}{n-\|i\|_1}$ if $n-d < \|i\| \leq n$ and $c_i = 0$ otherwise. If faults do occur then we recompute any u_i with $\|i\|_1 < n$ that was not successfully computed. If no faults occurred for any u_i with $\|i\| = n$ then we can again proceed with the classical combination. If faults affect any u_i with $\|i\|_1 = n$ then we add such i to the set J and then solve the GCP for $I_n \backslash J$. The solution is trivially obtained with hierarchical coefficients $w_i = 1$ for all $i \in I_n \backslash J$.

The largest solutions (in terms of unknowns) which may have to be recomputed are those with $\|i\|_1 = n-1$ which would be expected to take at most half the time of those solutions with $\|i\|_1 = n$. Since they take less time to compute they are also less likely to be lost due to failure. Additionally, there are $\binom{n-1+d-1}{d-1}$ solutions with $\|i\|_1 = n-1$ which is less than the $\binom{n+d-1}{d-1}$ with $\|i\|_1 = n$. As a result of these observations, we would expect to see far less disruptions caused by recomputations when using this approach compared to a lossless approach where all failed solutions are recomputed.

The worst case scenario with this approach is that all u_i with $\|i\|_1 = n$ are not successfully computed due to faults. In this case the resulting combination is simply a classical combination of level $n-1$. This only requires the solutions u_i with $n-d \leq \|i\| \leq n-1$. Likewise, all solutions to the GCP in this approach result in zero coefficients for all c_i with $i < n-d$. We can therefore reduce the overhead of the FTCT by only computing the solutions u_i for $n-d \leq \|i\|_1 \leq n$. It is known that the proportion of additional unknowns compared to the classical combination technique in this case is at most $\frac{1}{2(2^d-1)}$ [9].

The solutions u_i with $\|i\|_1 = n-1$ are only half the size of the largest u_i and hence recomputation of these may also be disruptive and undesirable. We could therefore consider recomputing only solutions with $\|i\|_1 \leq n-2$. By doing this the recomputations are even more manageable having at most one quarter the unknowns of the largest u_i. The worst case here is that all solutions with $i \geq n-1$ fail and we end up with a classical combination of level $n-2$. Again it turns out one does not require the entire downset I_n, in this case the (modified) FTCT requires solutions u_i with $n-d-1 \leq \|i\|_1 \leq n$. Using arguments similar to those in [9] it is easily shown that the overhead in this case is at most $\frac{3}{4(2^d-1)}$. The trade-off now is that the update of coefficients takes a little more work. We are back in the situation where we cannot guarantee that $(I_n \backslash J, \leq)$ is a lower semi-lattice.

To solve the GCP in this case we start with all w_i equal to 1. If failures affected any u_i with $\|i\|_1 = n$ we set the corresponding constraints $c_i = w_i = 0$. For failures occurring on u_i with $\|i\|_1 = n-1$ we have the constraints $w_i - \sum_{k=1}^{d} w_{i+e^k} = 0$ (with e^k being the multi-index with $e^k_l = \delta_{k,l}$). We note that (since the w_i are binary variables) this can only be satisfied if at most one of the w_{i+e^j} is equal to 1. Further, if $\sum_{k=1}^{d} w_{i+e^k} = 0$ we must also have $w_i = 0$. This gives us a total of $d+1$ feasible solutions to check for each such constraint. Given g failures on solutions with $\|i\|_1 = n-1$ we have at most $(d+1)^g$ feasible solutions to the GCP to check. This can be kept manageable if solutions are combined frequently enough that the number of failures g that are likely occur in between is small. One solves the GCP by computing the objective function (2.10) for each of the feasible solutions identified and selecting one which maximises this. Where some of the failures on the second layer are sufficiently far apart on the lattice, it is possible to significantly reduce the

number of cases to check as constraints can be optimised independently.

We could continue and describe an algorithm for only recomputing the fourth layer and below, however the coefficient updates here begin to become much more complex (both to describe and to compute). Our experience indicates that the recomputation of the third layer and below is a good trade-off between the need to recompute and the complexity of updating the coefficients. Our numerical results in Section 5 are obtained using this approach.

3.2. Probability of failure for computations. To analyse the expected outcome of the fault tolerant combination technique described in Section 3.1 we need to know the probability of each u_i failing. In particular, the availability of u_i will be modelled as a simple Bernoulli process U_i which is 0 if u_i was computed successfully and is 1 otherwise. It is assumed that each u_i is computed on a single computational node. Therefore we are interested in the probability that a failure occurs on this node before the computation of u_i is complete, that is $\Pr(U_i = 1)$. Suppose T is a random variable denoting the time to failure on a given node and the time required to compute u_i is given by t_i, then one has $\Pr(U_i = 1) = \Pr(T \leq t_i)$. One therefore needs to know something about the distribution of T.

Schroeder and Gibson analysed the occurrence of faults on 22 high performance machines at LANL from 1996-2005 [20]. They found that the distribution of time between failures for a typical node in the systems studied was best fit by the Weibull distribution with a shape parameter of 0.7. Based upon this study we will consider a model of faults on each node based upon the Weibull renewal process, that is a renewal process where inter-arrival times are Weibull distributed, with shape parameter $0 < \kappa \leq 1$.

There are several reasons for considering a renewal process for modelling faults. First, renewal theory is commonly used in availability analysis and there are many extensions such as alternating renewal processes in which one can also consider repair times. Second, we expect a fault tolerant implementation of MPI to enable the substitution of a failed node with another available node in which case computation can continue from some recovered state. This will be further discussed in Section 3.3. We now derive the value of $\Pr(U_i = 1)$.

Let $\{X_k\}_{k=1}^{\infty}$ be random variables for the successive times between failures on a node. We assume that the X_k are positive, independent and identically distributed with cumulative distribution

$$F(t) := \Pr(X_k \leq t) = 1 - e^{-(t/\lambda)^{\kappa}} \tag{3.1}$$

for some $0 < \lambda < \infty$ and $0 < \kappa \leq 1$. Let $S_k = \sum_{m=1}^{k} X_m$ (for $k \geq 1$) be the waiting time to the kth failure. Let $N(t)$ count the number of failures that have occured up until (and including) time t, that is $N(t) = \max\{k : S_k \leq t\}$). By the elementary renewal theorem one has

$$\lim_{t\to\infty} \frac{1}{t}\mathbb{E}[N(t)] = \frac{1}{\mathbb{E}[X_1]} = \frac{1}{\lambda\Gamma(1+\frac{1}{\kappa})}\,. \tag{3.2}$$

We thus get an expression for the (long term) average rate of faults.

Now, whilst we have a distribution for the time between failures, when a computation starts it is generally unknown how much time has elapsed since the last failure occurred. Hence our random variable T is what is referred to as the random incidence (or residual lifetime). Noting that one is more likely to intercept longer intervals of

the renewal process than shorter ones and that the probability distribution of the starting time is uniform over the interval, then it is straightforward to show [21] that T has density

$$g(t) = \frac{1 - F(t)}{\mathbb{E}[X_1]} = \frac{e^{-(t/\lambda)^\kappa}}{\lambda\Gamma(1 + \frac{1}{\kappa})}\,.$$

It follows that the cumulative probability distribution

$$G(t_i) := \Pr(T \le t_i) = \frac{1}{\lambda\Gamma(1 + \frac{1}{\kappa})} \int_0^{t_i} e^{-(x/\lambda)^\kappa} dx\,.$$

The resulting distribution has similar properties to the original Weibull distribution. In fact, when $\kappa = 1$ we note that X_k and T are identically distributed, they are exponential with mean λ. Further, for $0 < \kappa \le 1$ we have the following bound:

LEMMA 3.1. *For $0 < \kappa \le 1$ one has*

$$G(t) \le F(t)\,. \tag{3.3}$$

Proof. We note that via a change of variables that

$$\Gamma\left(1 + \frac{1}{\kappa}\right) = \int_0^\infty y^{\frac{1}{\kappa}} e^{-y} dy = \int_0^\infty \frac{\kappa}{\lambda}\left(\frac{x}{\lambda}\right)^\kappa e^{-(x/\lambda)^\kappa} dx$$

and therefore

$$\begin{aligned}
\Gamma\left(1 + \frac{1}{\kappa}\right) e^{-(t/\lambda)^\kappa} &= \int_0^\infty \frac{\kappa}{\lambda}\left(\frac{x}{\lambda}\right)^\kappa e^{-(x/\lambda)^\kappa} e^{-(t/\lambda)^\kappa} dx \\
&= \int_0^\infty \frac{x}{\lambda}\frac{\kappa}{\lambda}\left(\frac{x}{\lambda}\right)^{\kappa-1} e^{-(x/\lambda)^\kappa} e^{-(t/\lambda)^\kappa} dx \\
&= \left[-\frac{x}{\lambda} e^{-(x/\lambda)^\kappa} e^{-(t/\lambda)^\kappa}\right]_0^\infty - \int_0^\infty -\frac{1}{\lambda} e^{-(x/\lambda)^\kappa - (t/\lambda)^\kappa} dx \\
&= \int_0^\infty \frac{1}{\lambda} e^{-(x/\lambda)^\kappa - (t/\lambda)^\kappa} dx\,.
\end{aligned}$$

Since $0 < \kappa \le 1$ and $t, x \ge 0$ one has $x^\kappa + t^\kappa \ge (x + t)^\kappa$ and hence

$$\begin{aligned}
\Gamma\left(1 + \frac{1}{\kappa}\right) e^{-(t/\lambda)^\kappa} &\le \int_0^\infty \frac{1}{\lambda} e^{-((x+t)/\lambda)^\kappa} dx \\
&= \int_t^\infty \frac{1}{\lambda} e^{-(x/\lambda)^\kappa} dx \\
&= \int_0^\infty \frac{1}{\lambda} e^{-(x/\lambda)^\kappa} dx - \int_0^t \frac{1}{\lambda} e^{-(x/\lambda)^\kappa} dx \\
&= \Gamma\left(1 + \frac{1}{\kappa}\right) - \frac{1}{\lambda}\int_0^t e^{-(x/\lambda)^\kappa} dx\,.
\end{aligned}$$

Rearranging gives

$$\frac{1}{\lambda\Gamma(1 + \frac{1}{\kappa})} \int_0^t e^{-(x/\lambda)^\kappa} dx \le 1 - e^{-(t/\lambda)^\kappa}$$

which is the desired inequality. □

As a result of Lemma 3.1 one has that the probability of u_i failing to compute successfully is bounded above by

$$\Pr(U_i = 1) = G(t_i) \leq F(t_i)\,.$$

REMARK 2. *The result of Lemma 3.1 is essentially a consequence of the property*

$$\Pr(X \leq s + t \mid X > s) \leq \Pr(X \leq t) \tag{3.4}$$

where X is Weibull distributed with shape parameter $0 < \kappa \leq 1$ and $s, t \geq 0$. This property can be extended to the fact that if $s_2 \geq s_1 \geq 0$ then

$$\Pr(X \leq s_2 + t \mid X \geq s_2) \leq \Pr(X \leq s_1 + t \mid X \geq s_1)\,.$$

This has important implications on the order in which we compute successive solutions on a single node. Solutions one is least concerned about not completing due to a fault should be computed first and solutions for which we would like to minimise the chance of failure should be computed later.

REMARK 3. *We note that for $\kappa \geq 1$ the inequalities of Equations* (3.3) *and* (3.4) *are reversed thus $G(t) \geq F(t)$ and $\Pr(X \leq s + t \mid X > s) \geq \Pr(X \leq t)$ for $s, t \geq 0$.*

3.3. Fault Simulation in the FTCT algorithm. We first describe the parallel FTCT algorithm:

1. Given a (finite) set of multi-indices I, distribute the computation of component solutions u_i for $i \in I$ amongst the available nodes.
2. Each node begins to compute the u_i which have been assigned to it. For time evolving PDE's the solvers are evolved for some fixed simulation time t_s.
3. On each node, once a u_i is computed a checkpoint of the result is saved. If the node experiences a fault during the computation of a u_i, a fault tolerant implementation of MPI is used to replace this node with another (or continue once the interrupted node is rebooted). On the new node, checkpoints of previously computed u_i are loaded and we then assess whether the interrupted computation should be recomputed or discarded. If it is to be recomputed this is done before computing any of the remaining u_i allocated to the node.
4. Once all nodes have completed their computations they communicate which u_i have been successfully computed via a MPI_ALLREDUCE. All nodes now have a list of multi-indices $I' \subseteq I$ and can solve the GCP to obtain combination coefficients.
5. All nodes compute a partial sum $c_i u_i$ for the u_i that it has computed and then the sum is completed globally via a MPI_ALLREDUCE such that all nodes now have a copy of $u_{I'}^{\text{gcp}}$.
6. In the case of a time evolving PDE the u_i can be sampled from $u_{I'}^{\text{gcp}}$ and the computation further evolved by repeating from 2.

The optional step 6 for time evolution problems has many advantages. First, by combining component solutions several times throughout the computation one can improve the approximation to the true solution. Second, each combination can act like a global checkpoint such that when a u_i fails it can be easily restarted from the last combination (rather than the very beginning). Third, there are potential opportunities to reassess the load balancing after each combination and potentially re-distribute the u_i to improve performance in the next iteration.

For the numerical results in Section 5 we do not currently use a fault tolerant implementation of MPI and instead simulate faults by modifying the following steps:

2. Before each node begins computing the u_i assigned to it, a process on the node computes a (simulated) time of failure by sampling a distribution for time to failure and adding it to the current time. In our results we sample the Weibull distribution for some mean $\lambda > 0$ and shape $0 < \kappa \leq 1$.
3. Immediately after a u_i has been computed on a node we check to see if the current time has exceeded the (simulated) time of failure. If this is the case the most recent computation is discarded. We then pretend that the faulty node has been instantly replaced and continue with step 3 as described.

Note from Section 3.2 that sampling the Weibull distribution produces faults at least as often as the random incidence $G(t)$. Thus, by sampling the Weibull distribution in our simulation, the results should be no worse than what might occur in reality allowing for some small discrepancy in fitting the Weibull distribution to existing data of time between failures on nodes of a real machine.

The assumption in step 3 of replacing a failed node with another is based upon what one might expect from a fault tolerant MPI. In fact both Harness FT-MPI[1] and the relatively new ULFM specification[2] allow this, although it certainly does not occur in an instant as is assumed in our simulation. Due to limited data at the current time we are unable to predict what recovery times one might expect. We also note that (simulated) failures are checked for at the completion of the computation of each u_i. Since a failure is most likely to occur some time before the computation completes then time is wasted in the simulation from the sampled time of failure to the completion of the affected computation. Improving these aspects of the simulation and implementation with a fault tolerant MPI will be the subject of future work.

3.4. Expected error of the FTCT. In this section, we bound the expected interpolation error for the FTCT as applied to the classical combination technique as described in Section 3.1. In particular we look at the case where all solutions with $\|i\|_1 < n$ are recomputed, and the case where all solutions with $\|i\|_1 < n-1$ are recomputed as described in Section 3.1.

Given $u \in H^2_{\text{mix}}$, let

$$\epsilon_n := \frac{1}{3} \cdot 3^{-d} 2^{-2n} |u|_{H^2_{\text{mix}}} \sum_{k=0}^{d-1} \binom{n+d}{k} \left(\frac{1}{3}\right)^{d-1-k}$$

such that $\|u - u_n^c\|_2 \leq \epsilon_n$, see (2.3). Now given a (finite) set of multi-indices I we denote u_I^{gcp} to be a combination $\sum_{i \in I} c_i u_i$ which is a solution to the GCP described in Section 2.2. When faults prevent successful computation of some of the u_i we must find $u_{I'}^{\text{gcp}}$ for some $I' \subset I$. Consider the Bernoulli process $\{U_i\}_{i \in I}$ for which each $U_i = 0$ if u_i is computed successfully and is 1 otherwise as described in Section 3.2. Additionally it is assumed that the computation of each u_i is done within one node, that is many u_i can be computed simultaneously on different nodes but each individual u_i is computed within one hardware node. Let t_i be the time required to compute u_i for each $i \in I$. We assume that the time between failures on each node is Weibull distributed. As demonstrated in Section 3.2 the probability of each u_i being lost a the result of a fault is given by the random incidence distribution

$$\Pr(U_i = 1) = G(t_i) \leq F(t_i)\,.$$

[1] http://icl.cs.utk.edu/ftmpi/
[2] http://fault-tolerance.org/

Given that u_i with the same $\|i\|_1$ have a similar number of unknowns we assume they will take roughly the same amount of time to compute. We therefore define $t_k := \max_{\|i\|_1=k} t_i$, that is the maximal time to compute any u_i with level k.

As a result, for each $i \in I$ the probability of each u_i not completing due to failures is bounded by

$$\Pr(U_i = 1) \leq G(t_{\|i\|_1}) \leq F(t_{\|i\|_1}) \,.$$

. With this we can now give the main result.

PROPOSITION 3.2. *Given $d, n > 0$ and $I_n := \{i \in \mathbb{N}^d : \|i\|_1 \leq n\}$ let u_i be the interpolant of $u \in H^2_{mix}$ for $i \in I_n$. Let each u_i be computed on a different node of a parallel computer for which the time between failures on every node is independent and identically Weibull distributed with mean $\lambda > 0$ and shape parameter $0 < \kappa \leq 1$. Let t_i be the (wall) time required to compute each u_i and $t_n = \max_{\|i\|_1=n} t_i$. Suppose we recompute any u_i with $\|i\|_1 < n$ which is interrupted by a fault, let $\mathcal{I}$ be the set of all possible $I' \subseteq I_n$ for which u_i was successfully computed (eventually) iff $i \in I'$. Let $u^{gcp}_{\mathcal{I}}$ be the function-valued random variable corresponding to the result of the* FTCT *(i.e. $u^{gcp}_{I'}$ for some random $I' \in \mathcal{I}$), then the expected error is bounded above by*

$$\mathbb{E}\left[\|u - u^{gcp}_{\mathcal{I}}\|_2\right] \leq \epsilon_n \left(1 + 3\left(1 - e^{-(t_n/\lambda)^\kappa}\right)\right) .$$

Proof. Since u_i with $\|i\|_1 < n$ are recomputed we have that $U_i = 0$ for all $\|i\|_1 < n$ and therefore $\Pr(\mathcal{I} = I') = 0$ for $I_{n-1} \nsubseteq I'$. Note that I_k is a downset for $k \geq 0$, that is $I_k = I_k \downarrow$. Since the i with $\|i\|_1 = n$ are covering elements for I_{n-1} it follows that each of the I' for which $\Pr(\mathcal{I} = I') > 0$ are also downsets. It follows that there is a unique solution to the GCP for such I', namely $w_i = 1$ for all $i \in I'$ and $w_i = 0$ otherwise. That is $w_i = 0$ iff $U_i = 1$ and hence $w_i = 1 - U_i$. It follows that the error is bounded by

$$\|u - u^{\mathrm{gcp}}_{I'}\|_2 \leq \|u - u^c_n\|_2 + \sum_{\|i\|_1=n} |1 - w_i| \, \|h_i\|_2 = \|u - u^c_n\|_2 + \sum_{\|i\|_1=n} U_i \|h_i\|_2 \,.$$

From Lemma 3.1, the probability of a fault occurring during the computation of any u_i with $\|i\|_1 = n$ is bounded by $G(t_n)$ and therefore for $\|i\|_1 = n$ one has

$$\mathbb{E}[U_i] = 0 \cdot \Pr(U_i = 0) + 1 \cdot \Pr(U_i = 1) = \Pr(U_i = 1) = G(t_i) \leq G(t_n) \,.$$

It follows that

$$\begin{aligned} \mathbb{E}\left[\|u - u^{\mathrm{gcp}}_{\mathcal{I}}\|_2\right] &\leq \mathbb{E}\left[\|u - u^c_n\|_2 + \sum_{\|i\|_1=n} U_i \|h_i\|_2\right] \\ &= \|u - u^c_n\|_2 + \sum_{\|i\|_1=n} \mathbb{E}[U_i] \|h_i\|_2 \\ &\leq \|u - u^c_n\|_2 + \sum_{\|i\|_1=n} G(t_i) \|h_i\|_2 \,, \end{aligned} \tag{3.5}$$

and substituting the estimate $\|h_i\|_2 \leq 3^{-d}2^{-2\|i\|_1}|u|_{H^2_{\text{mix}}}$ yields

$$\begin{aligned}\mathbb{E}\left[\|u - u_{\mathcal{I}}^{\text{gcp}}\|_2\right] &\leq \epsilon_n + \sum_{\|i\|_1=n} G(t_n)3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}} \\ &\leq \epsilon_n + F(t_n)\sum_{\|i\|_1=n} 3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}} \\ &= \epsilon_n + \left(1 - e^{-(t_n/\lambda)^\kappa}\right)\binom{n+d-1}{d-1}3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}}\,.\end{aligned}$$

Now, noting that $\binom{n+d-1}{d-1} \leq \sum_{k=0}^{d-1}\binom{n+d}{k}(1/3)^{d-1-k}$ one has

$$\binom{n+d-1}{d-1}3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}} \leq 3\epsilon_n$$

and therefore,

$$\mathbb{E}\left[\|u - u_{\mathcal{I}}^{\text{gcp}}\|_2\right] \leq \epsilon_n\left(1 + 3\left(1 - e^{-(t_n/\lambda)^\kappa}\right)\right)\,. \tag{3.6}$$

Note that as $t_n/\lambda \to \infty$ we have $\mathbb{E}\left[\|u - u_{\mathcal{I}}^{\text{gcp}}\|_2\right] \leq 4\epsilon_n$. However, the worst case scenario is when $I' = I_{n-1}$ which results in a classical combination of level $n-1$ which has the error bound

$$\begin{aligned}\|u - u_{n-1}^c\|_2 \leq \epsilon_{n-1} &= \frac{1}{3}\cdot 3^{-d}2^{-2(n-1)}|u|_{H^2_{\text{mix}}}\sum_{k=0}^{d-1}\binom{n-1+d}{k}\left(\frac{1}{3}\right)^{d-1-k} \\ &\leq \frac{4}{3}\cdot 3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}}\sum_{k=0}^{d-1}\binom{n+d}{k}\left(\frac{1}{3}\right)^{d-1-k} \\ &= 4\cdot\epsilon_n\,.\end{aligned}$$

This is consistent with the upper bound (3.6) which is the desired result. ∎

Note the assumption that each u_i be computed on a different node is not necessary as we have bounded the probability of a failure during the computation of u_i to be independent of the starting time. As a result our bound is independent of the number of nodes that are used during the computation and how the u_i are distributed among them as long as each individual u_i is not distributed across multiple nodes.

The nice thing about this result is that the bound on the expected error is simply a multiple of the error bound for u_n^c, i.e. the result in the absence of faults. If the bound on $\|u - u_n^c\|_2$ was tight then one might expect

$$\mathbb{E}\left[\|u - u_{\mathcal{I}}^{\text{gcp}}\|_2\right] \lessapprox \|u - u_n^c\|_2\left(1 + 3\left(1 - e^{-(t_n/\lambda)^\kappa}\right)\right)\,.$$

Also note that (3.5) can be expressed as

$$\mathbb{E}\left[\|u - u_{\mathcal{I}}^{\text{gcp}}\|_2\right] \leq \|u - u_n^c\|_2 + \Pr(T \leq t_n)\|u_{n-1}^c - u_n^c\|_2\,.$$

If the combination technique converges for u then $\|u_{n-1}^c - u_n^c\|_2 \to 0$ as $n \to \infty$. Since $\Pr(T \leq t_n) \leq 1$ the error due to faults diminishes as $n \to \infty$. We now prove an analogous result for the case where only solutions with $\|i\|_1 < n-1$ are recomputed.

PROPOSITION 3.3. *Given $d, n > 0$ and $I_n := \{i \in \mathbb{N}^d : \|i\|_1 \leq n\}$ let u_i, t_i and t_n be as described in Proposition 3.2 with each u_i computed on different nodes*

for which time between failures is iid having Weibull distribution with $\lambda > 0$ *and* $0 < \kappa \leq 1$. *Additionally let* $t_{n-1} = \max_{\|i\|_1=n-1} t_i$. *Suppose we recompute any* u_i *with* $\|i\|_1 < n-1$ *which is interrupted by a fault, let* $\mathcal{I}$ *be the set of all possible* $I' \subseteq I_n$ *for which* u_i *was successfully computed (eventually) iff* $i \in I'$. *Let* $u_{\mathcal{I}}^{gcp}$ *be the function-valued random variable corresponding to the result of the* FTCT, *then*

$$\mathbb{E}\left[\|u - u_{\mathcal{I}}^{gcp}\|_2\right] \leq \epsilon_n \cdot \min\left\{16, 1 + 3\left(d + 5 - e^{-\left(\frac{t_n}{\lambda}\right)^\kappa} - (d+4)e^{-\left(\frac{t_{n-1}}{\lambda}\right)^\kappa}\right)\right\}.$$

Proof. This is much the same as the proof of Proposition 3.2. The solution to the GCP for I_n satisfies the property that if $w_i = 0$ for $\|i\|_1 = n-1$ then u_i was not computed successfully, that is $U_i = 1$. However the converse does not hold in general. Regardless, if $U_i = 1$ for $\|i\|_1 = n-1$ then the worst case is that $w_i = 0$ and $w_j = 0$ for the d possible $\|j\|_1 = n$ satisfying $j > i$. We therefore note that the error generated by faults affecting u_i with $\|i\|_1 = n-1$ is bounded by

$$\sum_{j \in I_n,\, j \geq i} \|h_j\|_2 \leq (d+4)3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}}. \tag{3.7}$$

Therefore we have

$$\begin{aligned}
\mathbb{E}\left[\|u - u_{\mathcal{I}}^{\text{gcp}}\|_2\right] &\leq \|u - u_n^c\|_2 + \sum_{\|i\|_1=n} G(t_n)\|h_i\|_2 \\
&\quad + \sum_{\|i\|_1=n-1} G(t_{n-1}) \sum_{j \in I,\, j \geq i} \|h_j\|_2 \\
&\leq \epsilon_n \left(1 + 3\left(1 - e^{-(t_n/\lambda)^\kappa}\right)\right) \\
&\quad + \binom{n-1+d-1}{d-1}\left(1 - e^{-(t_{n-1}/\lambda)^\kappa}\right)(d+4)3^{-d}2^{-2n}|u|_{H^2_{\text{mix}}} \\
&\leq \epsilon_n \left(1 + 3\left(1 - e^{-(t_n/\lambda)^\kappa}\right) + 3(d+4)\left(1 - e^{-(t_{n-1}/\lambda)^\kappa}\right)\right).
\end{aligned}$$

Now the expected error should be no more than the worse case which is $I' = I_{n-2}$ for which we have $\|u - u_{n-2}^c\|_2 \leq 16\epsilon_n$. Taking the minimum of the two estimates yields the desired result. ∎

To illustrate how this result may be used in practice, suppose we compute a level 12 interpolation in 3 dimensions on a machine whose mean time to failure can be modelled by the Weibull distribution with a mean of 100 seconds and shape parameter 0.7. Further, suppose u_i with $\|i\|_1 > 10$ are not recomputed if lost as a result of a fault and that t_{12} is estimated to be 1.0 seconds and t_{11} is at most 0.5 seconds. The expected error for our computation is bounded above by 1.63 times the error bound if no faults were to occur.

Whilst this provides some theoretical validation of our approach, in practice we can numerically compute an improved estimate by enumerating all possible outcomes and the probability of each occurring. The reason for this is that Equation (3.7) is an overestimate in general, particularly for relatively small d. In practice, a fault on u_i with $\|i\|_1 = n-1$ will generally result in the loss of $d-1$ of the largest hierarchical spaces in which case Equation (3.7) overestimates by a factor of $\frac{d+4}{d-1}$.

3.5. Expected computation time. We now repeat the above analysis, this time focusing on the mean time required for recomputations. The first issue to consider is that a failure may occur during a recomputation which will trigger another

recomputation. Given a solution u_i with $\|i\|_1 = m$, the probability of having to recompute r times is bounded by $G(t_m)^r \leq F(t_m)^r$. Hence the expected number of recomputations for such a u_i is bounded by

$$\sum_{r=1}^{\infty} rF(t_m)^r = \frac{F(t_m)}{(1-F(t_m))^2} = e^{(t_m/\lambda)^\kappa}(e^{(t_m/\lambda)^\kappa} - 1)\,.$$

Let the time required to compute each u_i be bounded by $t_i \leq c2^{\|i\|_1}$ for some fixed $c > 0$ and all $\|i\|_1 \leq n$. For a given $m \leq n$, suppose we intend to recompute all u_i with $\|i\|_1 \leq m$ upon failure, then the expected time required for recomputations is bounded by

$$R_m \leq \sum_{\|i\|_1 \leq m} t_{\|i\|_1} \sum_{r=1}^{\infty} rF(t_{\|i\|_1})^r \leq \sum_{k=0}^{m} \binom{k+d-1}{d-1} c2^k e^{(c2^k/\lambda)^\kappa}\left(e^{(c2^k/\lambda)^\kappa} - 1\right)\,.$$

and by bounding components of the sum with the case $k = m$ one obtains

$$\begin{aligned} R_m &\leq \binom{m+d-1}{d-1} e^{(c2^m/\lambda)^\kappa}\left(e^{(c2^m/\lambda)^\kappa} - 1\right) \sum_{k=0}^{m} c2^k \\ &\leq \binom{m+d-1}{d-1} e^{(c2^m/\lambda)^\kappa}\left(e^{(c2^m/\lambda)^\kappa} - 1\right) c2^{m+1}\,. \end{aligned}$$

The time required to compute all u_i with $\|i\|_1 \leq n$ once is similarly bounded by

$$C_n = \sum_{\|i\|_1 \leq n} t_i \leq \sum_{k=0}^{n} \binom{k+d-1}{d-1} t_k \leq \binom{n+d-1}{d-1} c2^{n+1}$$

and hence $R_m/C_n \approx (m/n)^{d-1}2^{m-n}e^{(c2^m/\lambda)^\kappa}(e^{(c2^m/\lambda)^\kappa} - 1)$ estimates the expected proportion of extra time spent on recomputations. We would generally expect that $c2^m/\lambda << 1$ (we assume the time to compute level m grids is much less than the mean time to failure) and therefore this quantity is small. As an example, if we again consider a level 12 computation in 3 dimensions for which $t_m \leq 2^{-(m-12)}$ and the time to failure is Weibull distributed with mean 100 seconds with shape parameter 0.7, the expected proportion of time spent recomputing solutions level 10 or smaller is $(10/12)^2 2^{-2} e^{(0.25/100)^{0.7}}(e^{(0.25/100)^{0.7}} - 1) \approx 2.68 \times 10^{-3}$. In comparison, if any of the u_i which fail were to be recomputed then a proportion of 4.23×10^{-2} additional time for would be expected for recomputations, almost 16 times more. Whilst this is a somewhat crude estimate it clearly demonstrates the our approach will scale better than a traditional checkpoint restart when faults are relatively frequent.

4. Implementation and Scalability. At the heart of our implementation is a very simple procedure: solve the problem on different grids, combine the solutions, and repeat. This section is broken up into different sub-sections based upon where different layers of parallelism can be implemented. We conclude by discussing some bottlenecks in the current implementation.

4.1. Top Layer: Load Balancing and Combination. The top layer is written in Python. As in many other applications, we use Python to glue together the different components of our implementation as well as providing some high level functions for performing the combination. It is based upon the development of NuMRF [17, 18]:

intended to be a clean interface where different computation codes can be easily interchanged or added to the application. This layer can be further broken down into 4 main parts. The first is the loading of all dependencies including various Python and *numpy* modules as well as any shared libraries that will be used to solve the given problem. In particular, bottom layer components which have been compiled into shared libraries from various languages (primarily C++ with C wrappers in our case) are loaded into Python using *ctypes*. The second part is the initialisation of data structures and construction/allocation of arrays which will hold the relevant data. This is achieved using PyGraFT [17, 18] which is a general class of grids and fields that allows us to handle data from the various components in a generic way. Also in this part of the code is the building of a sparse grid data structure. This is done with our own C++ implementation which was loaded into Python in the first part of the code. The third part consists of solving the given problem. This is broken into several "combination steps". A "combination step" consists of a series of time steps of the underlying solver for each component solution, followed by a combination of the component solutions into a sparse grid solution, and finally a sampling of the component solutions from the sparse grid solution before repeating the procedure. The fourth and final part of the code involves checking the error of the computed solutions, reporting of various log data and finalise/cleanup.

The top layer is primarily responsible for the coarsest grain parallelism, that is distributing the computation of different component solutions across different processes. This is achieved through MPI using *mpi4py*[3]. There are two main tasks the top layer must perform in order to effectively handle this. The first is to determine an appropriate load balancing of the different component solutions across a given number of processes. For our simple problem this can be done statically on startup before initialising any data structures. For more complex problems this can be done dynamically by evaluating the load balancing and redistributing if necessary at start of each combination step based upon timings performed in the last combination step. Re-distribution of the grids may require reallocating many of the data structures. The second task the top layer is responsible for is the communication between MPI processes during the combination step. This is achieved using two ALL_REDUCE calls. The first call is to establish which solutions have been successfully computed. This is required so that all processes are able to compute the correct combination coefficients. Each process then does a partial sum of the component solutions it has computed. The second ALL_REDUCE call then completes the combination of all component solutions distributing the result to all processes. Following this the component solutions are then sampled from the complete sparse grid solution.

4.2. Bottom Layer: Solver and Sparse Grid Algorithms. The bottom layer is made up of several different components, many of which are specific to the problem that is intended to be solved. When solving our advection problem we have 2 main components, one is responsible for the sparse grid data structure and functions relating to the sparse grid (e.g. interpolation and sampling of component solutions) and the other component is the advection solver itself. Both the sparse grid and advection solver components use OpenMP to achieve a fine grain level of parallelism. This is primarily achieved by distributing the work of large for loops within the code across different threads. The for loops have roughly constant time per iteration so the distribution of work amongst threads is done statically.

[3]http://mpi4py.scipy.org/

4.3. Optional Middle Layer: Domain Decompositions. The middle layer is currently being developed into the programming model. It is intended solely to handle various aspects relating to the computation of component solutions where domain decompositions are added as a third layer of parallelism. This will be achieved through an interface with a distributed array class of the NuMRF/PyGraFT framework at the top layer. This layer will need to interface with solver kernel from the bottom layer and then perform communication of data across domain boundaries. The combination of solutions onto the sparse grid in the top layer will also need to interface with this layer to handle the communication of different domains between MPI processes. It is intended that most of this will be transparent to the user.

4.4. Scalability bottlenecks. Since interpolation of the sparse grid and the solver (and any other time consuming operations) each benefit from load balancing with MPI and work sharing with OpenMP, any major hurdles to scalability will be caused by the ALL_REDUCE communication and any serial operations in the code (e.g. initialisation routines). Ignoring initialisation parts of the code it becomes clear we need to either reduce the size of the data which is communicated, or reduce the frequency at which it is communicated. The first can be done if we apply some compression to the data before communicating, i.e. we trade-off smaller communications for additional CPU cycles. Another way is to recognise that it is possible to do the ALL_REDUCE on a sparse grid of level $n-1$ if a partial hierarchisation is done to the largest component grids. This doesn't improve the rate in which the complexity grows but can at least reduce it by a constant. Reducing the frequency of the ALL_REDUCE can be done by performing partial combinations in place of full combinations for some proportion of the steps. This trades off the time taken to combine with some accuracy of the approximation. A partial combination is where a grid combines only with its neighbouring grids. However, the only way to really address the bottleneck caused by communication is to perform a full hierarchisation of the component grids [15]. By doing this one can significantly reduce the communication volume at the expense of increasing the number of messages. One can then reduce the number of messages by identifying those which are communicated to the same MPI processes. We currently have a first implementation of this which we intend to improve as development continues.

5. Numerical Results. In this section, we present some numerical results which validate our approach. The problem used to test our algorithm is the scalar advection equation

$$\frac{\partial u}{\partial t} + a \cdot \nabla u = 0$$

on the domain $[0,1]^3 \subset \mathbb{R}^3$ for constant $a \in \mathbb{R}^3$. For the results presented in this section we use $a = (1,1,1)$, periodic boundary conditions and the initial condition

$$u_0(x) = \sin(4\pi x_1)\sin(2\pi x_2)\sin(2\pi x_3)\,.$$

The PDE is solved using a Lax-Wendroff finite difference scheme giving results which are second order in space and time. We compare numerical solutions against the exact solution

$$u(x,t) = \sin(4\pi(x_1 - a_1 t))\sin(2\pi(x_2 - a_2 t))\sin(2\pi(x_3 - a_3 t))$$

to determine the solution error at the end of each computation.

TABLE 5.1
Numerical results for $r = 100$ runs for each l, τ using the Weibull distribution with mean of 1000 seconds and shape parameter of 0.7 for the fault simulation. The computation was performed on 2 nodes with 6 OpenMP threads on each.

l	τ	f_{ave}	ϵ_{ave}	ϵ_{min}	ϵ_{max}	w_{ave}	w_{min}	w_{max}
18	4	0.13	2.103e-4	2.096e-4	2.298e-4	26.96	26.94	27.07
20	5	0.37	7.064e-5	7.004e-5	8.580e-5	72.03	71.82	72.72
19	4	0.48	6.355e-5	6.266e-5	6.899e-5	131.1	130.8	131.9
21	5	1.21	1.979e-5	1.886e-5	3.030e-5	379.9	379.3	381.9
20	4	1.80	1.925e-5	1.856e-5	2.156e-5	649.8	648.2	653.3

TABLE 5.2
Numerical results for $r = 200$ runs for each l, τ using the Weibull distribution with mean of 1000 seconds and shape parameter of 0.7 for the fault simulation. The computation was performed on 6 nodes with 6 OpenMP threads on each.

l	τ	f_{ave}	ϵ_{ave}	ϵ_{min}	ϵ_{max}	w_{ave}	w_{min}	w_{max}
18	4	0.305	2.119e-4	2.096e-4	2.418e-4	9.216	9.196	9.392
20	5	0.535	7.166e-5	7.004e-5	1.170e-4	24.77	24.67	25.22
19	4	0.690	6.385e-5	6.244e-5	7.392e-5	44.55	44.40	45.90
21	5	1.805	2.015e-5	1.886e-5	3.179e-5	131.7	130.9	134.6
20	4	2.475	1.961e-5	1.844e-5	2.478e-5	224.4	223.3	228.7

A truncated combination technique as in (2.6) is used for our experiments. In order to apply the FTCT we need to compute some additional grids. We define

$$I_{n,\tau} := \{i \in \mathbb{N}^d : \min(i) \geq \tau \text{ and } n - d - 1 \leq \|i\|_1 \leq n\}$$

which is the set of indices for which we are required to compute solutions u_i if the top two levels are not to be recomputed in the event of a fault.

Note that as the grid sizes vary between the u_i so does the maximum stable time step size as determined by the CFL condition. We choose the same time step size for all component solutions to avoid instability that may otherwise arise from the extrapolation of time stepping errors during the combination. As a result our timesteps must satisfy the CFL condition for all component grids. By choosing Δt such that it satisfies the CFL condition for the numerical solution of $u_{(n-2\tau,n-2\tau,n-2\tau)}$ it follows that the CFL condition is also satisfied for all u_i with $i \in I_{n,\tau}$.

All of our computations were performed on a Fujitsu PRIMERGY cluster consisting of 36 nodes each with 2 Intel Xeon X5670 CPUs (6 core, 2.934GHz) with Infiniband interconnect.

5.1. Solution Error. We first looked at the effect of simulated faults on the error the computed solution. Given level n and truncation parameter τ the code was executed for some number of runs r on a fixed number of nodes using the same number of threads. Component solutions are combined twice in each run, once halfway through the time steps and again at the end. For each run we recorded the number of faults f that occurred, the l_1 error of the solution ϵ and the wall time w spent in the solver. We then calculated the average number of faults

$$f_{ave} = \frac{1}{r} \sum_{k=1}^{r} f_k \,,$$

TABLE 5.3

Numerical results for $r = 100$ runs for each l, τ using the exponential distribution with mean of 1000 seconds for the fault simulation. The computationa were performed on 2 nodes with 6 OpenMP threads on each.

l	τ	f_{ave}	ϵ_{ave}	ϵ_{min}	ϵ_{max}	w_{ave}	w_{min}	w_{max}
18	4	0.06	2.098e-4	2.096e-4	2.230e-4	26.96	26.93	27.10
20	5	0.17	7.098e-5	7.006e-5	1.155e-4	72.01	71.87	72.47
19	4	0.26	6.321e-5	6.266e-5	7.283e-5	131.1	130.8	131.7
21	5	0.71	1.942e-5	1.886e-5	2.980e-5	379.8	379.3	380.8
20	4	1.27	1.921e-5	1.856e-5	2.127e-5	649.3	647.9	653.4

TABLE 5.4

Numerical results for $r = 100$ runs for each l, τ using the exponential distribution with mean of 1000 seconds for the fault simulation. The computations were performed on 6 nodes with 6 OpenMP threads on each.

l	τ	f_{ave}	ϵ_{ave}	ϵ_{min}	ϵ_{max}	w_{ave}	w_{min}	w_{max}
18	4	0.070	2.103e-4	2.096e-4	2.296e-4	9.231	9.214	9.340
20	5	0.155	7.034e-5	7.004e-5	8.162e-5	24.80	24.70	30.83
19	4	0.265	6.327e-5	6.244e-5	7.039e-5	44.48	44.36	45.07
21	5	0.865	1.936e-5	1.886e-5	3.203e-5	131.7	131.0	134.5
20	4	1.415	1.921e-5	1.844e-5	2.178e-5	224.0	222.9	227.4

the average, minimal and maximal observed errors

$$\epsilon_{ave} = \frac{1}{r}\sum_{k=1}^{r} \epsilon_k \,, \quad \epsilon_{min} = \min\{\epsilon_1, \ldots, \epsilon_r\} \,, \quad \epsilon_{max} = \max\{\epsilon_1, \ldots, \epsilon_r\} \,,$$

and the average, minimal and maximal observed wall times

$$w_{ave} = \frac{1}{r}\sum_{k=1}^{r} w_k \,, \quad w_{min} = \min\{w_1, \ldots, w_r\} \,, \quad w_{max} = \max\{w_1, \ldots, w_r\} \,.$$

Table 5.1 shows our results for $r = 100$ runs of the FTCT with fault simulation on 2 nodes with 6 OpenMP threads on each. Faults were simulated as described in Section 3.3 using the Weibull distribution with mean of 1000 seconds and shape parameter 0.7 to sample the time between failures. As we increase the level n (or decrease τ) we increase the problem size and hence computation time. This in turn leads to an increase in the average number of faults that occur per run as seen in the f_{ave} column. The minimal error is the same as the error without failure (sometimes it is fractionally smaller). Comparing with the average error we see that the additional error generated by recovery from simulated faults is small. Also worth noting is that the variability in computation time is quite small indicating that any recomputations, when they occur, do not seem to cause any significant disruptions.

In Table 5.2 we repeat this experiment with $r = 200$ runs on 6 nodes with 6 OpenMP threads on each. Whilst running with additional nodes leads to a decrease in computation time we experience more faults on average because of the additional nodes. However, we can see that the effect of the increased average number of faults is quite small on both the average solution error and the average wall time.

Table 5.3 again shows results for $r = 100$ runs of the FTCT with fault simulation on 2 nodes with 6 OpenMP threads on each. However, for this experiment the faults

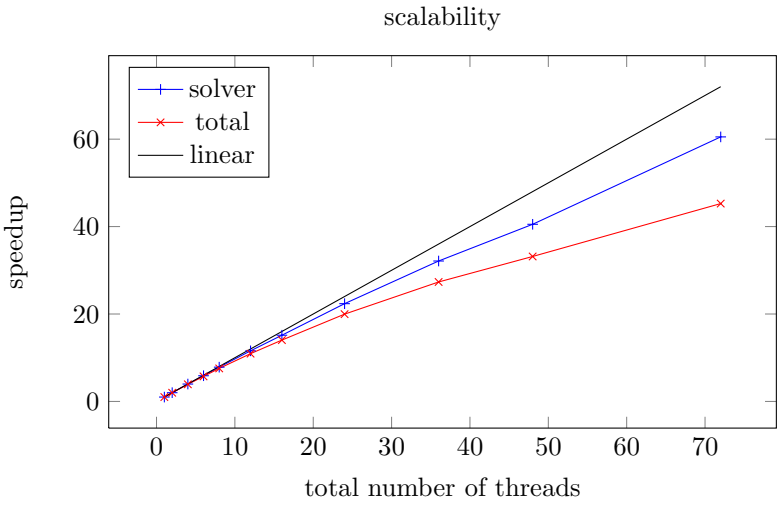

FIG. 5.1. *This plot demonstrates the scalability of our implementation for the advection problem using a* $l = 22, \tau = 6$ *truncated combination. Fault simulation was disabled and only one combination was performed at the end of all computations.*

are exponentially distributed with a mean of 1000 seconds. We see that for this distribution the faults are a little less frequent on average leading to a slightly smaller average error. Similar is observed in Table 5.4 where we repeat the experiment with $r = 200$ runs on 6 nodes with 6 OpenMP threads on each. Here the average number of faults is substantially less than the results of Table 5.2 and this is again reflected by a smaller average error in comparison. The large w_{max} in the 2nd row is due to a single outlier, the next largest time being 25.34. No simulated faults occurred for this outlier so we suspect it was due to a system issue.

5.2. Scalability. In Figure 5.1 we demonstrate the scalability and efficiency of our implementation when the fault simulation is disabled. Noting from Table 5.2 that faults have very little effect on the computation time we expect similar results with fault simulation turned on. The advection problem was solved using a $n = 22, \tau = 6$ truncated combination. The component solutions were combined only once at the end of the computation. The *solver* time reported here is the timing of the core of the code, that is the repeated computation, combination and communication of the solution which is highly scalable. The *total* time reported here includes the time for Python to load modules and shared libraries, memory allocation and error checking. The error checking included in the *total* time is currently computed in serial and could benefit from OpenMP parallelism.

In Figure 5.2 we compare the computation time required for our approach to reach a solution compared to more traditional checkpointing approaches, in particular, with a local and global checkpointing approach. With global checkpointing we keep a copy of the last combined solution. If a failure affects any of the component grids it is assumed that the entire application is killed and computations must be restarted from the most recent combined solution. We emulate this by checking for faults at each combination step and restart from the last combination step if any faults have occurred. With local checkpointing each MPI process saves a copy of each component solution it computes. In this case when a faults affect component solutions we need

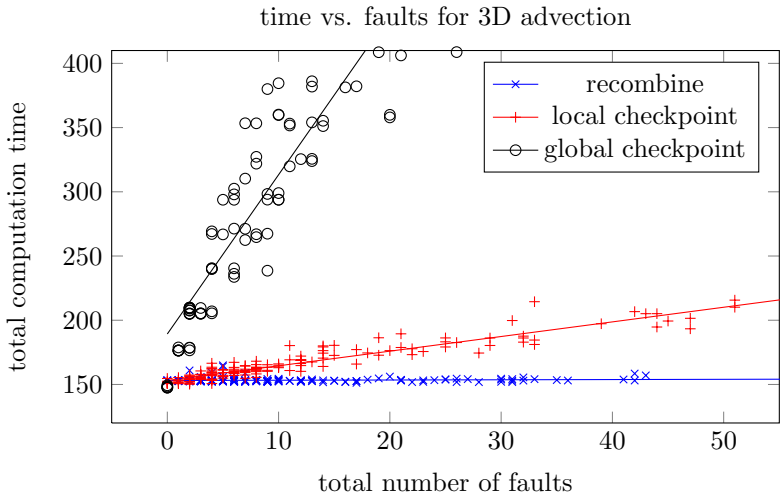

FIG. 5.2. *We compare the time taken to compute the solution to the 3D advection problem using three different approaches to fault tolerance. The problem size is fixed at level 21 with truncation parameter 5. All computations used 6 MPI processes with 6 OpenMP threads each. Component solutions are combined 4 times throughout the computation and it is during the combination that we check for faults. The* recombine *method is our approach described in Section 3.1. The* local checkpoint *method involves each* MPI *process checkpointing component solutions. The* global checkpoint *method involves all* MPI *processes checkpointing the last combined solution. For each method the problem was run numerous times with* MTTF *varying from* 25 *to* 1000 *seconds.*

only recompute the affected component solutions from their saved state. In both checkpointing methods the extra component solutions used in our approach are not required and are hence not computed. As a result these approaches are slightly faster when no faults occur. However, as the number of faults increases, it can be seen from Figure 5.2 that the computation time for the local and global checkpointing methods begins to grow. A line of best fit has been added to the figure which makes it clear that the time for recovery with global checkpointing increases rapidly with the number of faults. Local checkpointing is a significant improvement on this but still shows some growth. On the other hand our approach is barely affected by the number of faults and beats both the local and global checkpointing approaches after only a few faults. For much larger number of faults our approach is significantly better.

Conclusion. A generalisation of the sparse grid combination technique has been presented. From this generalisation a fault tolerant combination technique has been proposed which significantly reduces recovery times at the expense of some upfront overhead and reduced solution accuracy. Theoretical bounds on the expected error and numerical experiments show that the reduction in solution accuracy is very small. The numerical experiments also demonstrate that the upfront overheads become negligible compared to the costs of recovery using checkpoint-restart techniques if several faults occur. There are some challenges associated with load balancing and efficient communication with the implementation of the combination technique. Studying these aspects and improving the overall scalability of the initial implementation will be the subject of future work. As the ULFM specification continues to develop, the validation of the FTCT on a system with real faults is also being investigated.

REFERENCES

[1] G. Bosilca, R. Delmas, J. Dongarra, and J. Langou, *Algorithm-based fault tolerance applied to high performance computing*, in J. Parallel Distrib. Comput., 69 (2009), pp. 410–416.

[2] F. Cappello, A. Geist, B. Gropp, L. Kale, B. Kramer, and M. Snir, *Toward exascale resilience*, in Int. J. High Perform. Comput. Appl., 23(4) (2009), pp. 374–388.

[3] F. Cappello, *Fault tolerance in petascale/exascale systems: current knowledge, challenges and research opportunities*, in Int. J. High Perform. Comput. Appl., 23(3) (2009), pp. 212–226.

[4] J. Dean, and S. Ghemawat, *MapReduce: simplified data processing on large clusters*, in Commun. ACM, 51(1) (2008), pp. 107–113.

[5] J. Garcke, *Sparse grids in a nutshell*, in In Jochen Garcke and Michael Griebel, editors, Sparse Grids and Applications, Lecture Notes in Computational Science and Engineering, 88 (2013), pp. 57–80.

[6] G. Gibson, B. Schroeder, and J. Digney, *Failure tolerance in petascale computers*, in CT-Watch Quarterly, 3(4) (2007), pp. 4–10.

[7] M. Griebel, and H. J. Bungartz, *Sparse grids*, Acta Numerica, 13 (2004), pp. 147–269.

[8] M. Griebel et. al., *A combination technique for the solution of sparse grid problems*, in Iterative Methods in Linear Algebra (Brussels, 1991), pp. 263–281.

[9] B. Harding and M. Hegland, *A robust combination technique*, in S. McCue, T. Moroney, D. Mallet, and J. Bunder, editors, Proceedings of the 16th Biennial Computational Techniques and Applications Conference, ANZIAM Journal, 54 (CTAC2012), pp. C394–C411.

[10] B. Harding and M. Hegland, *Robust solutions to PDEs with multiple grids*, Sparse Grids and Applications - Munich 2012, J. Garcke, D. Pflüger (eds.), Lecture Notes in Computational Science and Engineering 97, Springer, 2014, to appear.

[11] B. Harding and M. Hegland, *A parallel fault tolerant combination technique*, M. Bader, A. Bode, H.-J. Bungartz, M. Gerndt, G.R. Joubert, F. Peters (eds.), Parallel Computing: Accelerating Computational Science and Engineering (CSE), Advances in Parallel Computing 25, IOS Press, 2014, pp. 584–592.

[12] M. Hegland, *Adaptive sparse grids*, Proceedings of CTAC, ANZIAM Journal, 44 (2003), pp. C335–C353.

[13] M. Hegland, J. Garcke, and V. Challis, *The combination technique and some generalisations*, Linear Algebra and its Applications, 420 (2007), pp. 249–275.

[14] K. Huang and J. Abraham, *Algorithm-based fault tolerance for matrix operations*, in IEEE Trans. Comput., 33(6) (1984), pp. 518–528.

[15] P. Hupp, R. Jacob, M. Heene, D. Pflüger, and M. Hegland, *Global Communication Schemes for the Sparse Grid Combination Technique*, M. Bader, A. Bode, H.-J. Bungartz, M. Gerndt, G.R. Joubert, F. Peters (eds.), Parallel Computing: Accelerating Computational Science and Engineering (CSE), Advances in Parallel Computing 25, IOS Press, 2014, pp. 564–573.

[16] R. Karp, *Reducibility among combinatorial problems*, in Complexity of Computer Computations: Proc. of a Symp. on the Complexity of Computer Computations, R. E. Miller and J. W. Thatcher, Eds., The IBM Research Symposia Series, New York, NY: Plenum Press, 1972, pp. 85–103.

[17] J. Larson, M. Hegland, B. Harding, S. Roberts, L. Stals, A. Rendell, P. Strazdins, M. Ali, C. Kowitz, R. Nobes, J. Southern, N. Wilson, M. Li, Y. Oishi, *Fault-Tolerant Grid-Based Solvers: Combining Concepts from Sparse Grids and MapReduce*, Proceedings of 2013 International Conference on Computer Science (ICCS), Procedia Computer Science, Elsevier, 2013.

[18] J. Larson, M. Hegland, B. Harding, S. Roberts, L. Stals, A. Rendell, P. Strazdins, M. Ali, J. Southern, *Managing complexity in the parallel sparse grid combination technique*, M. Bader, A. Bode, H.-J. Bungartz, M. Gerndt, G.R. Joubert, F. Peters (eds.), Parallel Computing: Accelerating Computational Science and Engineering (CSE), Advances in Parallel Computing 25, IOS Press, 2014, pp. 593–602.

[19] A. Murarasu, J. Weidendorfer, G. Buse, D. Butnaru, D. Pflüger, *Compact data structure and scalable algorithms for the sparse grid technique*, ACM SIGPLAN Notices, 46(8), ACM 2011.

[20] B. Schroeder and G. Gibson, *A large-scale study of failures in high-performance computing systems*, in Dependable Systems and Networks, International Conference on, 2006, pp. 249–258.

[21] K. Trivedi, *Probability and Statistics with Reliability Queuing and Computer Science Applications*, John Wiley & Sons, Inc., New York, 2002.